

Canada's AI imperative
Overcoming risks, building trust

omnia**AI**



Former Governor General of Canada David Johnston wrote in his book *Trust* that, “Adhering to the moral imperative ahead of the operational imperative builds and maintains trust.” At Deloitte, we firmly believe Canada can and should lead the way in artificial intelligence. To do so, we need leaders who will address AI's risks and ethical implications in order to foster greater trust in it and capitalize on its full potential. We need leaders like **Sir Frederick Banting**, the Canadian medical practitioner who won the 1923 Nobel Prize in Medicine for co-discovering insulin, who is depicted on the cover of this report.

Table of contents

- Introduction 2
- AI is different and so are the risks..... 4
- What we heard from Canadians 8
- AI needs more than a technical fix—it requires building trust 20
- Canada's imperative to lead..... 32
- Definitions and endnotes 34

Artificial intelligence (AI) is expected to be one of the leading economic drivers of our time, and Canada has the opportunity—and responsibility—to be a global leader. As a country, we have the research strength, talent pool, and startups to capitalize, but that's not enough if we want to lead in an AI-driven world and shape what it might look like. True leadership is required—that means taking steps now to establish a world-class AI ecosystem in Canada.

Introduction

AI has the potential to be the catalyst for an era of unprecedented innovation, progress, and prosperity. Yet many worry that the rapid pace of AI development and adoption is outstripping our ability to understand and manage AI's impact on society and the economy.

The increasing use of AI has brought a host of new risks and challenges into the spotlight. Around the world, citizens are raising questions about AI's impact on privacy, security, bias, consumer protection, and more. And they're looking to business and government leaders to provide substantial and satisfactory answers and solutions to those questions.


We believe the questions and concerns raised by Canadians and others around the world are rooted in two key areas: understanding and trust. People struggle to understand what AI is—and isn't—and what it can offer. It doesn't help that the hype surrounding AI on all sides can make it difficult to separate science from science fiction. This blurry understanding is in turn contributing to a rising distrust of AI among businesses, governments, and citizens alike.

Left unaddressed, this lack of trust could have a serious impact on Canada's future prosperity. Businesses, governments, and other organizations could be persuaded to slow or even abandon AI investments. Governments could implement regulations that severely limit the use of AI technologies. Such steps might assuage the worries of the public, but at what cost to our competitiveness in the years to come?

At Deloitte, we believe that Canada is uniquely positioned to respond to the challenges of AI. It's time for our nation to step to the forefront of this global conversation and take the lead. It's time for us to work together to build AI standards and practices that mirror our distinctly Canadian values, and develop AI that is open, safe, equitable, and used in ways that create prosperity for all.

We must begin this effort by first understanding the perceptions Canadians have about AI, so that business and government can address those attitudes most effectively. This report, the second installment in a multi-part Deloitte series on Canada's AI imperative, sheds light on Canadians' key perceptions about AI—and offers our recommendations for moving forward.

We believe the questions and concerns raised by Canadians and others around the world are rooted in two key areas: **understanding and trust.**



AI is different and so are the risks

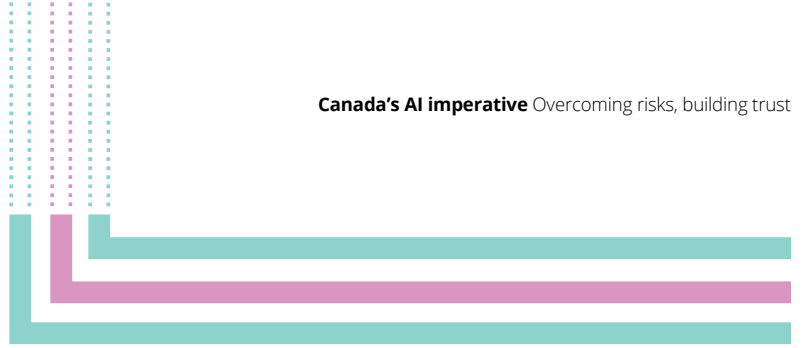
It can be tempting to dismiss the concerns about the risks of AI as similar to any other new technology.

Certainly, common questions related to cost, technical feasibility, and change management all hold true. But a growing body of research shows that AI *is* in fact different than other technologies, and this gives greater importance to AI's risks and ethical dimensions.¹

- **The speed and scale** at which AI is expected to proliferate and combine with existing technologies shortens the amount of time available to develop responses to new risks as they emerge.
- **The unknown reach and impact** of AI—given its potential effect on almost all industries and sectors simultaneously—makes planning for every possible risk seem daunting.

Our research shows that the ethical risks of AI are top of mind for both businesses and consumers. Between July and September 2018, Deloitte surveyed over 1,000 Canadian citizens and 2,500 businesses from around the world to understand the state of AI in Canada.²


- Canadian companies' reservations about AI often involved AI's "black box" problem. The lack of explanation for and unintended consequences of AI decisions, and the use of AI to manipulate information and create falsehoods, emerged as top concerns.
- Among Canadian consumers, 65 percent reported privacy concerns over how companies were using their data today. Understanding is a major issue as well: only 4 percent of Canadians felt confident about explaining what AI is and how it works.³



Given the rapid pace of change with respect to AI, it's not surprising that businesses and consumers are wary and confused. The potential pervasiveness of AI and its applications does open the door to myriad risks (see *What are the risks of AI?*). From a technical perspective, concerns about algorithmic bias, data privacy, and cybersecurity are real challenges to overcome. As well, large companies' heavily reported missteps have only contributed to a public perception that AI can be highly risky and harmful—a perception that matters, since a lack of understanding of AI will have a real effect on organizational and public policy decision-making.⁴

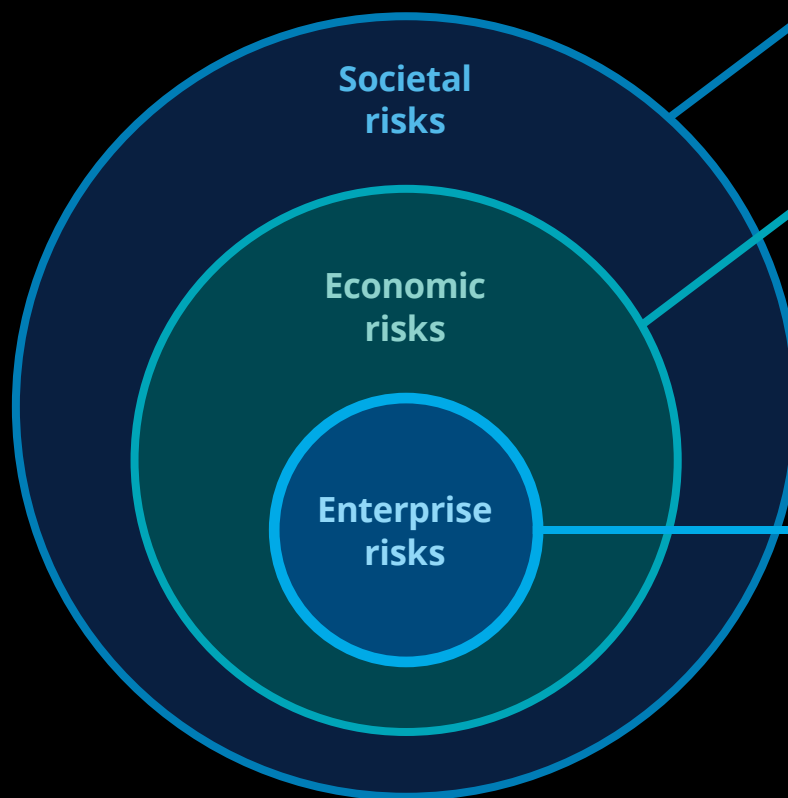
To ensure Canada reaps the benefits of AI in the years to come, it's vital that Canada's businesses and governments work together to help Canadians better understand AI—and develop the controls, rules, and practices necessary to address legitimate concerns about the impact of AI on our society. To help shed some light on the issue, in the pages that follow we explore what our research tells us about Canadians' current perceptions about AI.

From a technical perspective, concerns about algorithmic bias, data privacy, and cybersecurity are real challenges to overcome.



What are the risks of AI?

The speed and scale at which AI is evolving makes it harder to identify risks in an effective and timely manner. They could also be manifesting themselves in unfamiliar ways. We broadly classify these risks based on the varying impacts each one will have on society, markets, and enterprise. These risks are not mutually exclusive—they are all interconnected in some way.



Societal risks





The way that AI interacts with citizens can pose risks from a societal perspective. Consider, for instance, the use of automated surveillance which could infringe on individual and cultural views of privacy and autonomy, or the use of automated drones for military warfare.⁵

Economic risks

The transformational nature of AI can cause significant changes to current economic and market structures. For instance, the changing nature of jobs and the workforce can potentially alter employment opportunities for individuals and organizations.

Enterprise risks

Enterprise risks arise as an organization looks to incorporate AI into its strategic and technical frameworks.⁶ These risks might include:

-  **Privacy:** AI creates the potential for personal data to be misused or compromised.
-  **Transparency:** The use of black-box models can limit the ability to explain decisions or audit outcomes.
-  **Control:** AI creates the potential for AI systems or models to produce unintended consequences.
-  **Cybersecurity:** AI systems, models, and data may be vulnerable to cyber threats and adversarial attacks.

What we heard from Canadians

While discussions about AI risks and ethics are increasingly common in business and policy circles, you need to speak directly with Canadians to truly understand the sources of distrust and confusion surrounding AI.

Our in-depth interviews uncovered five themes that illuminate Canadians' experiences with AI and can help guide business and government as they move forward.

The fears and concerns that were raised don't have easy answers. There are numerous contradictions and tensions in what we heard from our interview subjects (see *Exploring the contradictions*). As expected, we heard about long-term fears surrounding the potential impact of artificial general intelligence (AGI) alongside concerns about how narrow applications of AI, like computer vision or machine learning, are being used today and could be used tomorrow.

At Deloitte, we generally distinguish between two different types of ethical challenges when it comes to AI governance: ethical breaches and ethical dilemmas. In our interviews, concerns over high-profile ethical breaches (such as the Facebook-Cambridge Analytica scandal) quickly segued into concerns that, as a society, we're ill prepared to address the long-term ethical dilemmas raised by AI.⁷

Ethical breaches are violations of trust that occur when there is an established and accepted code by which to determine if wrongdoing has occurred (e.g., codes of ethics, laws, regulations, clear norms). Many of the issues raised by AI about privacy, security, and transparency fall into this category. Pre-existing norms, business practices, and laws are a starting point, but leaders need to figure out how to apply and update them with regard to AI. Ethical breaches will require much discussion and iteration, and can likely be addressed with technical, legal, or strategic fixes.

Ethical dilemmas occur when the question of which moral code should be applied is contested (e.g., if it must hit a person, should the autonomous car hit the child or the elderly person?). In these cases, the way forward is unclear and not everyone agrees as many moral frameworks could apply. Hence, the dilemma. Questions related to AI's purpose and values—where and how it should be used—fall into this category, as do questions about the future nature of collaboration between AI and humans in the workplace. These are longer-term challenges that will likely require new forms of collaboration and discussion to address them. At the same time, this isn't the first time that technology has hit up against profound questions. Consider the profession of bioethics as a practical, applied model. What we don't yet have is as rich a body of AI ethics as we have with bioethics.

The narratives that follow do not explain all concerns about AI or speak for all Canadians. Instead, they're meant to complement existing research and highlight some of the nuanced barriers that must be addressed to promote greater trust and adoption.

Our methodology

We went coast to coast across Canada conducting interviews with Canadians to understand how they think and feel about AI with the goal of translating these concerns into recommendations for policymakers and businesses.

This research was conducted using Doblin's method of design research, which is an applied research approach that borrows from ethnographic methods. Doblin is Deloitte's innovation and human-centered design consultancy. Human-centered design (HCD) helps organizations identify opportunities for change by first understanding people's behaviours and experiences. HCD leans on established methods from the social sciences that systematically study and develop insight into how people, cultures, and society function.

In total, we conducted 13 two-hour, in-home interviews and observational research to hear how people reflected on and rationalized their experiences and to see their lives in context. Such qualitative research allows us to understand both the facts and the meaning behind the facts (for instance, if only 4 percent of Canadians understand AI, why is that?), bringing depth and nuance to our analysis of people's perceptions of and experiences with AI. Our sampling size was chosen to gather the right level of depth of data over sheer quantity.

Interviewees were intentionally selected to represent a range of geographies, ages, income levels, educational attainment, and gender identities: variables from the Organisation for Economic Cooperation and Development (OECD) data that correlate to technological disruption. We also selected participants working in industries likely to be disrupted by technological change and those working in industries that were least likely to be disrupted, which were drawn from Statistics Canada NAICS codes and OECD data. The final 13 interviewees were selected through an extensive screening process.



Our interviewees told us that while AI feels ubiquitous, they aren't clear about when and where they are using it—fuelling a mythology around AI that's more science fiction than science.

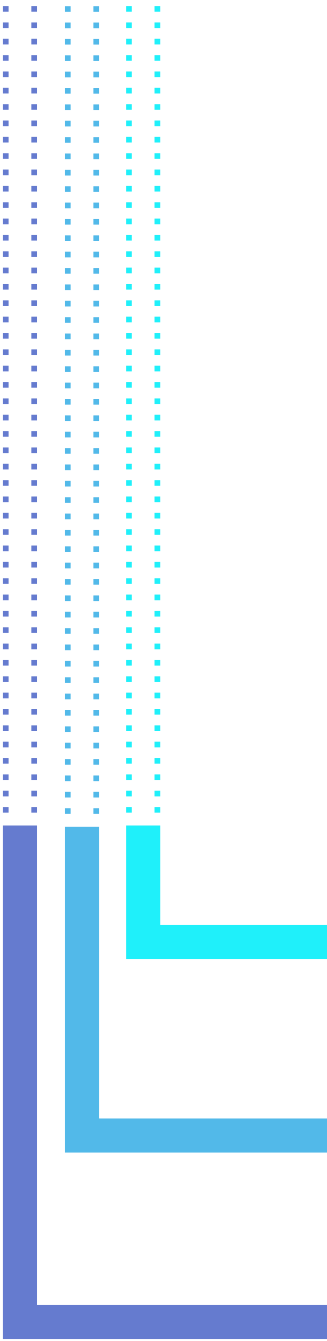
Our interviewees assumed they were constantly interacting with AI. Many expressed they couldn't imagine an AI-free life, as technology seems so pervasive. Despite this, they reported feeling deeply uninformed about how and when they were using AI. They often drew upon examples from popular culture, science fiction, or high-profile news stories. Some of our interviewees defined AI as algorithms while others spoke about robots.

The intangibility of AI may be contributing to a sense of confusion and distrust. This concern should not be dismissed. Knowledge of an innovation cannot be held in the hands of a select few. Research shows that knowledge and understanding are critical factors in successful technology diffusion across organizations and societies.⁸ And our ongoing research into barriers to greater AI adoption found that a lack of understanding among executives and employees, particularly the non-technical expertise required to apply AI to business challenges, is a challenge for Canadian companies in scaling their AI efforts beyond a proof of concept. Businesses and governments will need to work together to demystify AI and build AI literacy, which means explaining the basics of how AI works, the ways in which it uses data, and what it can and cannot do.

Exploring the contradictions:

What is AI truly capable of?

Most of our interviewees told us that AI would not be capable of doing their jobs. They were optimistic that AI would augment human intelligence, not replace it. On the other hand, some of our interviewees spoke of a world where AI surpassed human intelligence and was able to communicate in languages and make decisions beyond our comprehension.



"I'm not a techy person, so I would say no, I don't use AI in my day-to-day life. I know people who do, but I personally don't. Unless I do and I don't know, maybe?"

"In terms of my role, I would say that there's no AI that could take it on; my job is a very human connection-type job."

"I think for more rote or automated tasks, that's where AI adds a lot of value in terms of efficiency and reduction of human error."



"AI will treat us like ants."


"There will be areas where these algorithms, simply through learning over and over again, are gonna surpass our abilities."



Our interviewees expressed concerns that increased AI will lead to a loss of control and decision-making over their own lives. They emphasized the idea that, unlike decisions made by humans, there is no clear accountability when AI fails.

Worries over data privacy and security were common. Interviewees believed that opaque privacy policies often resulted in their data being collected or shared without permission, possibly being used to make decisions that could harm them or society at large. Other major concerns related to accountability. Interviewees expressed worry that AI would be able to make and impose unfair decisions in important situations, like job hiring or criminal convictions, without clear accountability or oversight. Much of the concern seemed focused on the idea that algorithms might make decisions on their own. This speaks to the idea of *algorithmic aversion*, a preference for human decision-makers even when an algorithm is proven to outperform humans.⁹

Left unaddressed, there's a risk that consumers and businesses may come to see AI as more of an autocrat than an advisor, and would avoid it. Businesses and governments will need to work together to combat this by considering use cases where humans should always be kept in the loop and work to increase transparency and predictability in terms of how AI is applied. Our interviewees also called for more clarity in how their data is used and pointed to governments to lead in creating accountability and safeguards for AI use.



Exploring the contradictions:

Do we value privacy or convenience?

Privacy concerns were a recurring theme in our interviews. Interviewees expressed distrust toward businesses and had low expectations of businesses keeping their data safe. Many expressed fatigue from worrying about data privacy. Nevertheless, they readily shared data with private companies in order to have services tailored to their individual needs and interests.

"Privacy, not to say I don't value it, but I think it's already out there anyway so I can't be as concerned anymore."


"Businesses adopting stricter data privacy regulations. I think that's really important."

"I know some people feel very uncomfortable with how their data might be used unbeknownst to them. Obviously, if it were used for nefarious reasons, such as a political agenda, as it's been in the news, that's not good if it's used as a device to create division. On the other hand, I kind of like it when I'm suggested something that I might not have thought of."



"I don't feel like I'm a conspiracy theory person or anything like that, but I feel that you really just have to be maybe a bit wary of technology."

"So who's responsible at a point when a machine makes a decision and something happens? Is it...the technology? The company that made the technology? Or is it the person who chose to buy that technology or the person who's in charge of that project?"



Our interviewees were not only concerned about how AI could affect their lives now; they were worried about humanity's place in an AI-driven world.

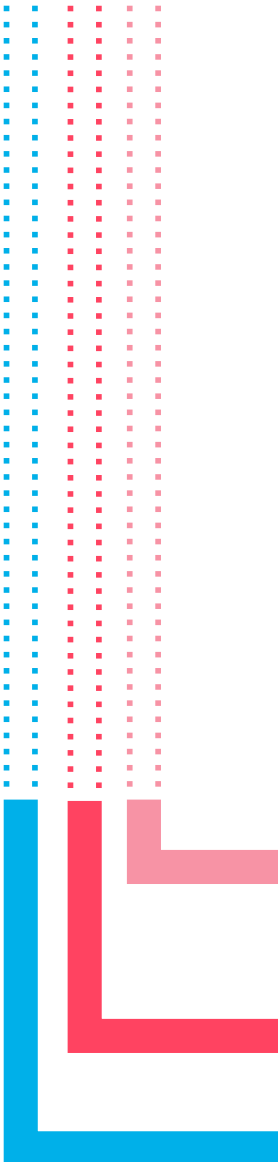
Speculation over what AI might be capable of made it easy for our interviewees to switch from their current worries to long-term fears about the relevance of human existence. They cited attributes such as empathy and emotional intelligence as uniquely human characteristics, but were uncertain if AI could surpass these human qualities in the future. More so than other technologies, artificial intelligence seems to touch a nerve when it comes to difficult questions about what constitutes our own humanity.

Fearing the loss of control and human agency will only fuel resistance and distrust of AI. Businesses and governments will need to work together to assuage these fears and create a strong shared vision of the ways in which AI will lead to prosperity for all. This includes giving the workforce of today—and tomorrow—the skills needed to thrive in a new, AI-driven world of work.

Exploring the contradictions:

Is AI human?

Our interviewees believed that AI must reflect humanity, both the good and the bad, as we are the ones who created it. At the same time, they were cognizant of the fact that AI is heavily data-dependent and therefore does not display emotions or qualities that are human-like.



"So, human beings don't have the edge on making the world a better place. So, maybe robots could do a better job. But then why do we even exist at all?"

"Empathy is a very human trait. You cannot build that into AI."

"I wonder about art. I wonder if it's worth defending as strictly human."

"I think one of the weaknesses of collecting human-type data is that we can't be figured out with a computer, and so huge mistakes are going to be made."

"Writing and literature and journalism are such human pleasures. Without them, what is there?"

"Knowledge is power. If you consider the self-checkout, I can make an informed decision not to use it because it's taking a job away from a person. I would rather have a person than a machine do that job."




Our interviewees told us that it's not about removing all bias from AI. They're worried that AI will replicate the worst of our biases and remove the good, like the human ability to exercise judgment, faith, and hope.


Our interviews demonstrated that recent news stories about AI making biased decisions, from gendered hiring to racialized pre-trial detention decisions, are dominating public sentiment. In fact, some interviewees felt strongly that AI development should stop until these errors are corrected. Yet others balked at what they saw as the opposite: AI making what they saw as purely rational decisions instead. They felt that a data-driven machine could never capture the best of the human spirit, for example, the ability to see promise in a young offender or to understand suffering when delivering a terminal medical diagnosis. Instead, they wanted to keep the best of humanity, but hoped technology might help us remove some of the worst.

Businesses and governments will need to work together to assure consumers and citizens of the fairness of AI decision-making or risk losing the social licence to operate. This means creating AI that purposefully amplifies the values we wish to see. The ethical development of AI will need to be front and centre from the beginning of development, and proper redress mechanisms should be put in place in the event that AI goes wrong after the fact.

Exploring the contradictions:
Is bias only an AI problem?

Interviewees expressed strong concerns about the inability of AI to think like a human when making decisions. The human capacity to judge situations with empathy and contextual awareness was mentioned often as key to making the right, or best, decisions. On the other hand, interviewees acknowledged that humans also have implicit biases that can lead to bad decision-making, biases that could possibly be avoided by using AI.





Our interviewees told us they worry about how ready we are as a society for the pace of change that AI will bring.

They reported feeling overwhelmed by the pace of technological change and felt unsure whom they could trust to lead in the face of widespread change caused by AI. We heard that people felt governments could be “too slow” or “ill-equipped” to react given the current pace of change. While businesses might be better suited to acting quickly, prominent cases of ethical breaches related to data privacy, bias, or malicious use of information were cited by our interviewees as examples of businesses not taking their obligations to lead seriously.

There will always be an inherent tension between pushing for technological progress and safeguarding those who are affected by it. By working together, businesses and governments can demonstrate to consumers and citizens that they take their concerns seriously and are working toward a future with AI prosperity for all.

Exploring the contradictions:
**Who should take
the lead on AI?**

Some of our interviewees expressed a desire for strong AI leadership from government but worried that government might move too slowly to be effective. Others felt that the ones designing the technology should be leading, but they were concerned those businesses would be driven by self-interest.

"I think that, ideally, businesses as the designers and developers of these technologies have obligations and self-interested obligations to be moral stewards of their tech."

"Right now, businesses aren't the right spokespeople for AI; they won't be believed."

"People think AI tech is benefiting them, but it's basically the companies that are winning."

"With AI, you're jumping from a horse-drawn carriage into a Ferrari overnight. I hope society knows how to build traffic lights."

"When new things happen, can government keep up with all the changes? They can't for health care, so I can't imagine they could for AI."

"Businesses don't speak for me. I don't pay my taxes to a corporation. For me, the only publicly accountable institution is government. That's why it exists."

AI needs more than a technical fix—it requires building trust

To take full advantage of AI's potential benefits, we need to build public trust—something sorely missing today. Canadians' feelings of distrust in AI go deep; they are rooted in emotion and can't be overcome by mere technical fixes.

Unfortunately, business and government may not have much time to act to address the perceived risks of AI before Canadians definitively turn against the new technology. Our research found that respondents struggled to distinguish between short- and long-term AI risks (see *Are short- and long-term risks a false distinction?*). No matter the technical issue at the heart of an AI-related breach—

poor data collection, a lack of oversight, etc.—respondents were quick to look at it in terms of big-picture ethical dilemmas surrounding AI's use.

If this distrust and misunderstanding of AI continues to grow, we risk seeing its adoption slow and backlash increase. A fatal self-driving car accident in Arizona, for instance, has resulted in more than 20 incidents of individuals attacking self-driving cars being tested.¹⁰

Left unchecked, this lack of trust will contribute to a societal divide between those who trust and use AI and those who don't.¹¹ Should this happen, it could hinder efforts to develop an AI prosperity strategy that benefits everyone, because the tools and solutions created won't address all the important issues or meet all the important needs.



For Canada to lead in addressing these AI challenges, we'll need collaboration between our policymakers, businesses, and academia as well as with international bodies. The time to do so is now, at the current stage of AI development, while the stakes are still relatively low. This is key to ensuring that Canada's governments and businesses become trusted leaders in shaping the future of AI.

In the pages that follow, Deloitte offers our recommendations for what actions Canada's business and government leaders should take to foster a greater understanding of and trust in AI. This is not meant to be an exhaustive list, but simply a starting point and a call to action for leaders to meaningfully address Canadians' real concerns about AI.

Are short- and long-term risks a false distinction?

There's research to show that, in fact, the tendency to classify AI risks as either short- or long-term issues might not be the right approach at all.¹² By classifying issues into two categories largely disconnected from one another, there's an impulse to continuously prioritize the short-term risks over the long-term ones. This focus may come at a price: the decisions we make in the future will be constrained by the ones we make today.¹³

Consider the development of the internet, which is probably the closest example we have to the proliferation of AI. It was built without a view toward longer-term security issues and we're still paying the price for that today.¹⁴ AI is similar in the sense that if we only focus on providing technical fixes for the problems we see now, we're unlikely to be prepared for the issues that arise in the future.

If this distrust and misunderstanding of AI continues to grow, we risk seeing its adoption slow and backlash increase.



Bring Canadian values to the forefront

Amid the confusion and uncertainty surrounding AI, there is a clear opportunity for Canada to step forward and provide global leadership. AI needs to be inclusive, fair, open, and safe, values that Canadians hold dear.

Key actions for business

Start with values and ethics. Before embarking on any AI project, consider the ethical implications of the tool you are building and codify how your organization will ensure you maintain fairness, privacy, safety, and openness. This is not an exercise that can be done after the fact; it must be the starting point. For instance, think about the project's fit with company values, its ability to deliver a safe and reliable solution, the societal impact of the technology, and potential reputational implications. Any application should comply with existing legislation and codes of practice, such as compliance with Canadian and organization-specific data security and privacy laws.

Create inclusive teams—and use inclusive data. Canadian businesses are uniquely placed to lead by drawing on the diverse, skilled talent pool available within our borders. Diverse and inclusive development teams are critical for building solutions that incorporate and meet the

needs of diverse audiences. Recent news of racialized facial-recognition systems serves as a sobering example of why it's important to get inclusion right.

Data is another critical source in preventing and addressing bias. Think about how data provenance affects the accuracy, reliability, and representativeness of the solution. Moreover, to prevent unintentional bias, consider whether there is accurate representation of all the groups you are trying to serve in the data you use. Bias should be viewed from a wide scope, including social, gender, and financial; indeed, any bias that might adversely affect disparate groups. Data audits and algorithmic assessments are important for flagging and mitigating discrimination and bias in machine-learning models. For the longer term, Canadian businesses can lead the way in developing robust data practices and standards that allow data to be shared freely and openly while ensuring it isn't used in ways that could be harmful.

Diverse and inclusive development teams are critical for building solutions that incorporate and meet the needs of diverse audiences.

 **Key actions for government**

Act as a vocal leader and convener on global standards. Influencing AI standards globally is a clear opportunity for government to both assuage the fears of Canadians at home and establish Canada as a leader on the global stage. The federal government has taken early steps to collaborate with industry groups like the CIO Strategy Council on this important work, but there's a need to move quickly.¹⁵ The race to set global standards is already heating up and having a voice at the table is critical for retaining a leadership role. Canada was the first country to announce a national AI strategy and now we need to maintain this momentum.



Ensure AI literacy is on the agenda

If AI continues to conjure more images of dystopian futures than ground-breaking innovations, then Canadians will continue to be distrustful. Business and government leaders need to work together to combat the hype around AI and build literacy instead. Canada should lead the way in ensuring its population is AI-fluent and ready to put this understanding to use. This means helping Canadians develop an understanding of how AI works and what it can and cannot do.



Key actions for business

Build non-technical fluency among employees and executives. Businesses can lead by investing in training employees to pair technical fluency with business understanding. This doesn't mean that employees and executives at your organization need to learn how to code algorithms. Quite the opposite. It means pairing an applied understanding of AI and its capabilities with existing non-technical skillsets so that the value of applying AI to solve business problems is clear.

This also means embedding ethics at every stage of the AI process. All employees and executives should take a holistic view in identifying and addressing ethical considerations throughout the AI development life cycle, not just at the beginning. For example, establishing a problem scope should also mean being mindful of the potential implications of solving it with AI. And, when deploying, monitoring, and improving AI solutions, employees and executives should ensure they are scaled in a way that respects the organization's internal values and operations.

Canada should lead the way in ensuring its population is AI-fluent and ready to put this understanding to use.

Key actions for government

Make AI literacy a priority. Canada should make AI literacy a pillar of its AI strategy. Our interviewees echoed this as well—they felt information about AI should be disseminated through a public information campaign, similar to the ones for cannabis and mental health. Governments around the world are starting to tackle this issue and Canada should, too. For example, Finland has launched an open online course aimed at getting 1 percent of the population trained in AI basics.¹⁶ These emerging practices could provide a starting point for Canadian governments at all levels to assess the best ways to build knowledge about AI. This could be by offering public education platforms similar to the Finnish model or by finding made-in-Canada ways to work AI knowledge into school curricula.



Collaborate to create the right controls and incentives

Businesses and governments need to work in tandem to build the necessary guardrails to assure Canadians they're working to prevent, detect, and repair AI breaches. We need a clear regulatory environment and avenues for redress, alongside sound organizational governance, to enable innovation and growth that don't compromise on trust.




Key actions for business

Proactively build oversight and accountability controls. Procedures for the development, maintenance, and oversight of AI systems need to be in place from the very beginning across the entire AI life cycle to detect and flag errors so that potential breaches can be detected before they occur and addressed promptly if they do. Managing AI risk and ethical concerns requires a focused approach, starting with understanding the role of AI in corporate strategy and the resultant

AI risks, and taking action to drive value-aligned AI innovation. Furthermore, building AI innovation will mean embedding values into the development life cycle, from use case definition, modelling, and training to deployment and monitoring.

Policies addressing matters such as when to put a human in the loop and when and how to conduct AI impact assessments will help prevent and mitigate any AI errors.



Clear and consistent regulation is essential for creating a stable business environment for innovators and building public trust.




Key actions for government

Ensure the regulatory environment promotes trust and innovation. Clear and consistent regulation is essential for creating a stable business environment for innovators and building public trust. Regulations must balance protection and privacy with the need for exploration and advancement. Tools like regulatory sandboxes can help governments keep pace with innovation and produce better long-term outcomes by providing a controlled environment to test products and ideas.¹⁷

This allows the government to understand the myriad regulatory changes required alongside industry before codifying them.

Create new legal frameworks for algorithmic accountability. In situations where people believe AI has wrongly or unfairly affected them, new legal frameworks that lay out clear redress and accountability for the consequences of decisions made by AI algorithms will need to be created.





Inform, engage, and communicate —continuously

Transparency can't be a one-off activity. It means finding the right balance between information overload and keeping people in the dark. Think of the exchange you have when you go to see your doctor about a procedure or diagnosis: you don't need to know every medical detail but you want to be adequately informed about the impact, the risks, and any preparations you need to make.



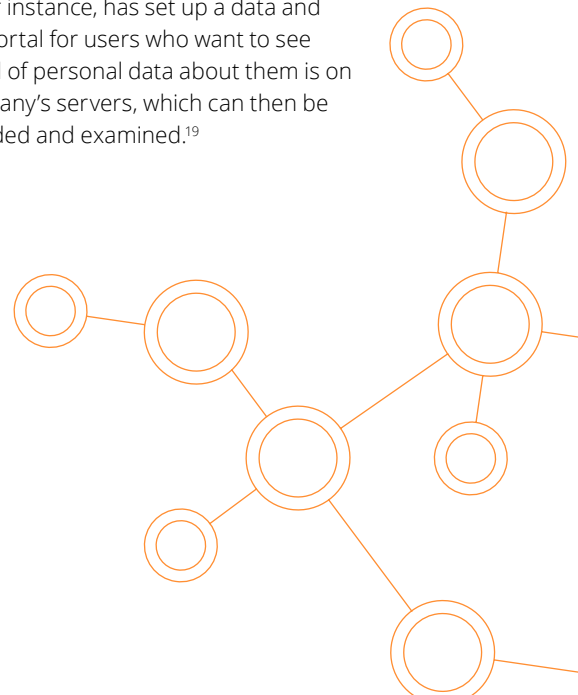
Key actions for business

Clearly articulate why AI is necessary.

Businesses can combat low trust in and understanding of AI by actively sharing the positive value that AI brings to their work. Sharing stories of successful AI projects with employees and, in some cases, the broader public is critical for getting employees and consumers on the same page. You must be able to articulate the value AI will bring, especially since the fear of shifting roles or job losses is high. Being clear in messaging and having strong champions articulate the value and expectations will be key for gaining acceptance. As an example, Scotiabank's stated objectives for interactive AI systems begin with a clear directive: those systems must be truly useful. They need to improve outcomes for customers and society as well as the bank.¹⁸

Make privacy an ongoing conversation.

Shift the burden away from the user by making privacy policies and user agreements easy to understand. Keep these agreements top of mind by offering frequent reminders or pop-ups that make clear to consumers how their data is being used and allow them to change their permissions. Apple, for instance, has set up a data and privacy portal for users who want to see what kind of personal data about them is on the company's servers, which can then be downloaded and examined.¹⁹





Being clear in messaging and having strong champions articulate the value and expectations will be key for gaining acceptance.




Key actions for government

Engage the public and guide the conversation. Making individual Canadians feel heard is critical for building trust. But the ability to engage citizens effectively hinges on having a well-informed public, which can be challenging when it comes to complex topics like AI. To give citizens a voice in the development of AI policy decisions, the government should consider experimenting with innovative consultation and engagement approaches such as citizen reference panels or the one being developed in Estonia (see *International spotlight*).²⁰

International spotlight: Estonia

Estonia recognizes the importance of public consultations on AI—and the difficulty of consulting a public that feels uninformed. To create a shared understanding of AI, the country's leadership has drawn on its rich cultural mythology to provide an impactful metaphor: the Kratt. The Kratt is a magical creature in Estonian lore made of hay or household items that can be brought to life. Creating this parallel has enabled the national AI taskforce to further public dialogue about the responsibilities and accountabilities of “the Kratt” and draft an algorithmic liability law—known as the Kratt law.²¹




Prepare for machine-human collaboration

As the use of AI in society increases, we will need to start reskilling and preparing the workforce to take up higher-value tasks and jobs for the long term. This includes thinking about how people and AI will work together rather than focusing on a people-versus-AI outcome.

Key actions for business

Retrain employees if AI is going to change their jobs. Businesses should begin investing in employees to prepare them for the changing nature of work. Over the long term, employees will require ongoing skills training and career counselling as new skills become in demand. One example is RBC's digital navigation program, which offers training to over 20,000 branch managers to be able to transition into positions that require similar skillsets.²²

Explore opportunities to redesign work to take advantage of AI and learning. Businesses will need to focus beyond automation and identify the most promising areas in which AI can augment workers' performance as they shift into more creative and value-added work. There needs to be a plan for redesigning and reinventing work that will combine the capabilities of machines and people, create meaningful jobs and careers, and help employees with the learning and support they will need to navigate rapidly evolving circumstances.²³




Businesses should begin investing in employees to prepare them for the changing nature of work.

 **Key actions for government**

Provide support for lifelong learning. As the half-life of skills rapidly shrinks, constant reskilling and retraining will need to become the norm. However, we also heard from interviewees the need for programs that will develop the right kind of multi-purpose skills, like empathy and critical thinking, to support lifelong learning.

Governments should consider developing shorter and more targeted certifications that are created in conjunction with businesses and incorporate co-op or apprenticeship terms to meet these needs. Consider, for instance, the way Singapore's government works with industry to assess the current state of workplace learning and offers guidance on effectively structuring workplace learning activities.²⁴



Canada's imperative to lead


AI has the potential to fundamentally change the way we live, work, do business, and make decisions in every aspect of our lives. It could unlock a new and exciting era of great shared prosperity for Canadians and people around the world. But we must first address the legitimate concerns about how AI will be developed and used before we can truly realize its immense rewards. Because at this point in time, the broader public is growing ever more skeptical—and ever more distrustful—of AI.


Trust in AI is falling because we are becoming increasingly aware of the risks, challenges, and ethical implications of using it, even as we grow more astonished at the opportunities it could unlock. Addressing these risks, overcoming these challenges, and tackling these ethical implications is essential if we are to foster greater trust in AI and capitalize on its full potential.

Canada's global reputation and commitment to transparency, inclusion, and collaboration presents our country—and our people—with a unique opportunity to lead the world in developing sound, moral, and ethical AI practices that put people first and deliver prosperity to the many, not just the few. As former Governor General of Canada David Johnston wrote in his book *Trust*, "Adhering to the moral imperative ahead of the operational imperative builds and maintains trust."

Working together, Canadian governments, businesses, and academia can address citizens' AI worries, helping them understand what AI is, what it isn't, and how it could benefit all of us. Collaboration will also be essential for ensuring we address AI's ethical dilemmas in a way that protects society yet allows for innovation, growth, and shared prosperity. "What is most important," writes Johnston, "and what builds trust, is making sure we get to our destination together."²⁵

If we truly join forces, that destination will be Canada's leadership position in shaping how AI is used in the future, both at home and around the world.





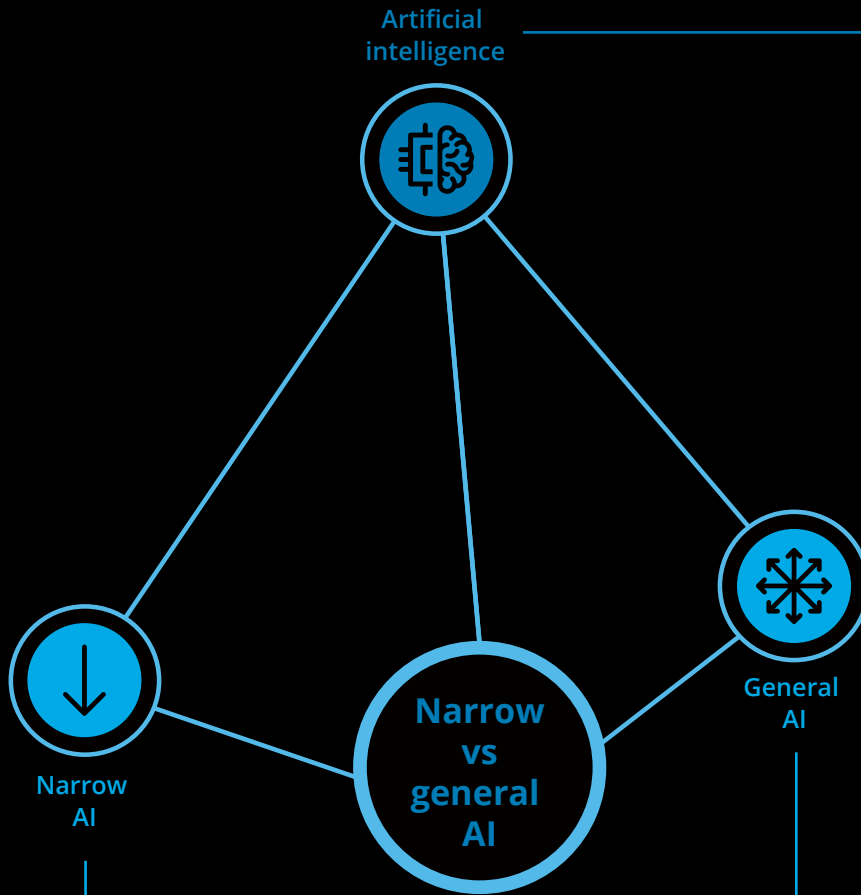
“Trust is earned incrementally and granted with greater acquaintance. The same cannot be said for the loss of trust. It often speeds off at the first sign of doubt and usually at the first suggestion of deceit.”

– David Johnston, Former Governor General of Canada

Definitions

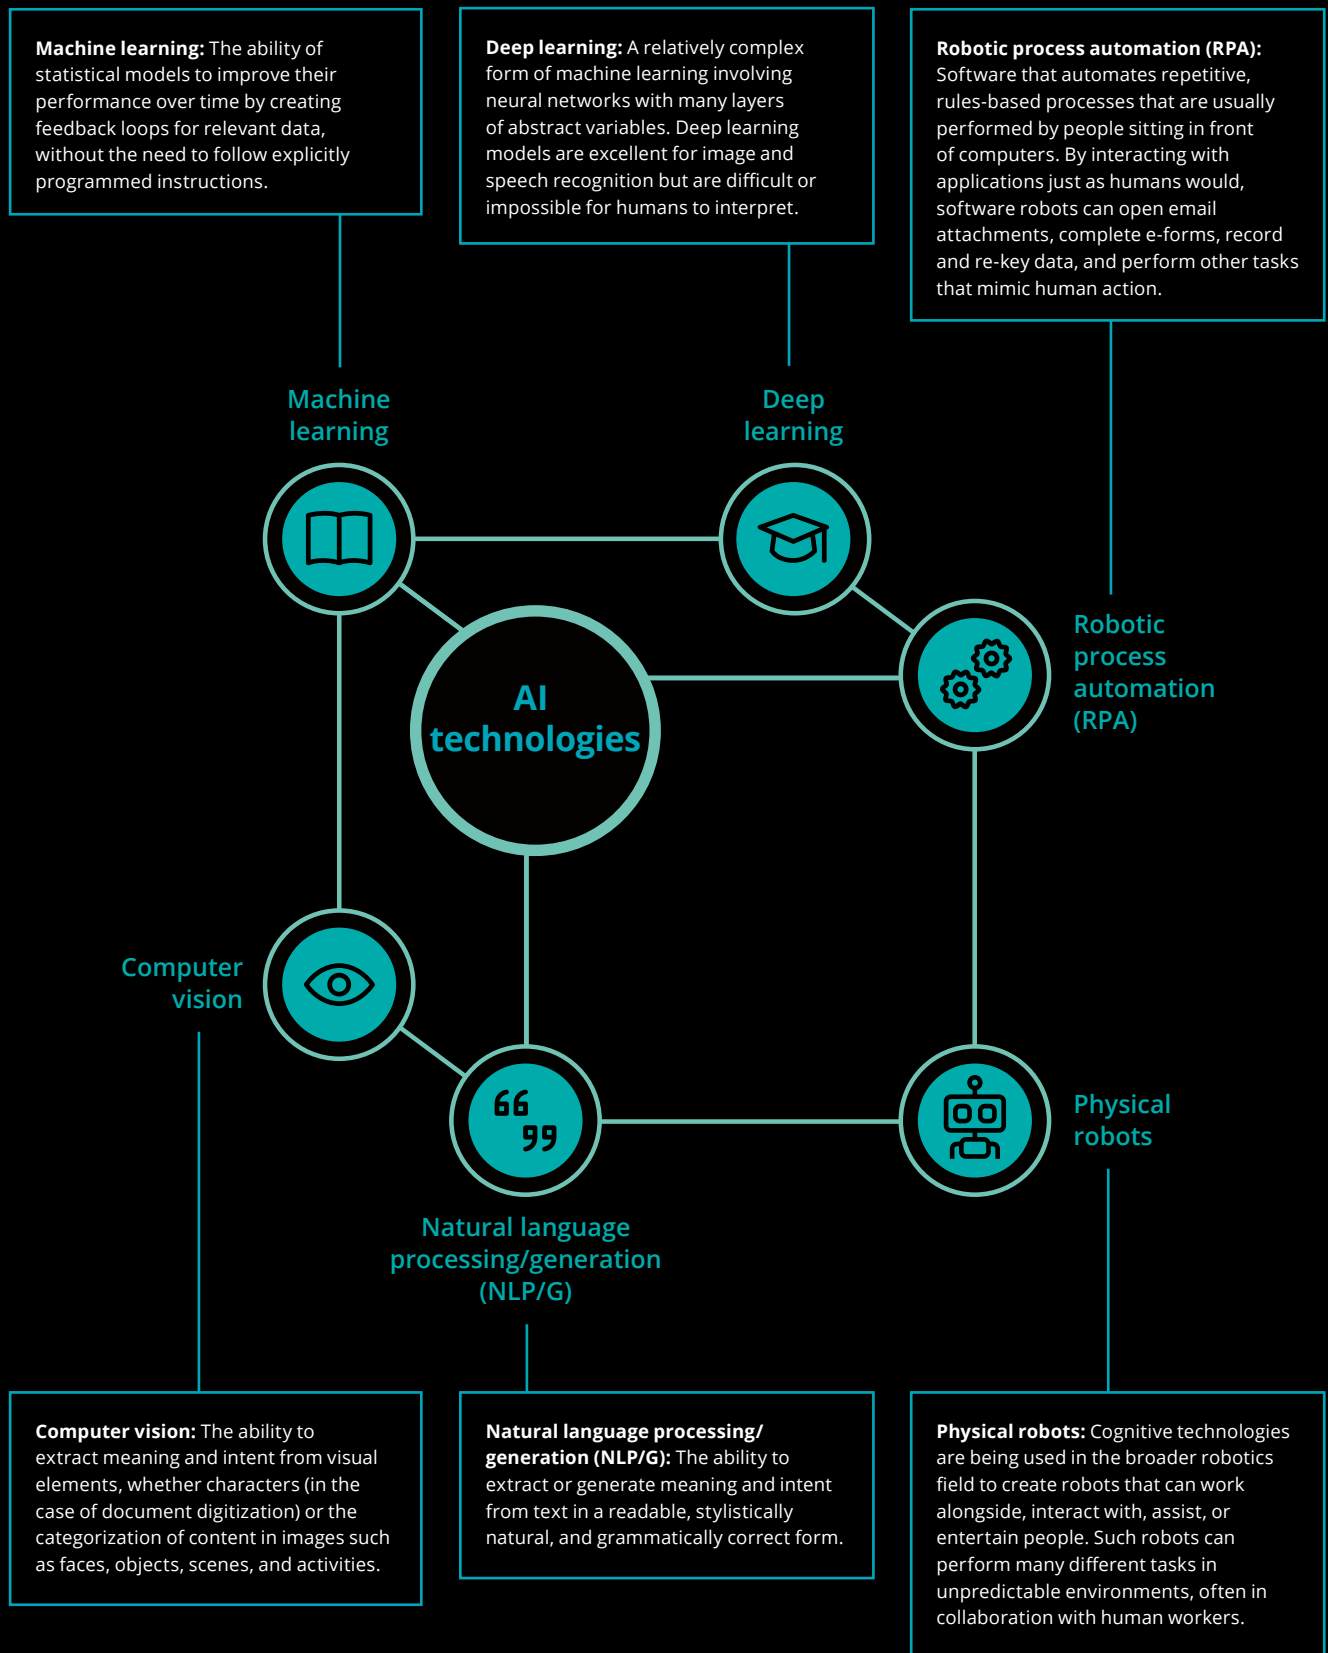
Artificial intelligence: Deloitte defines AI as systems and applications that perform tasks that mimic or augment human intelligence, ranging from simple gaming to sophisticated decision-making agents that adapt to their environment.

People often use the term AI interchangeably with the technologies that enable it—including machine learning, deep learning, neural networks, natural language processing, rule engines, robotic process automation, and combinations of these capabilities for higher-level applications.



Narrow AI: Most current applications are what we call narrow AI, which means it can only perform the task it was designed to do. This means that for every problem, a specific algorithm needs to be designed to solve it.⁴

General AI: The holy grail of AI is general AI, a single system that's equal or superior to human intelligence. Today's AI applications don't exhibit this 'human-like' intelligence—an algorithm designed to drive a car, for example, wouldn't be able to detect a face in a crowd or guide a domestic robot assistant.⁵



Endnotes

1. Iain Cockburn, Rebecca Henderson, and Scott Stern, "The Impact of Artificial Intelligence on Innovation" (Cambridge, MA: National Bureau of Economic Research, March 2018), <https://doi.org/10.3386/w24449>, Ajay Agrawal, John McHale, and Alex Oettl, "Finding Needles in Haystacks: Artificial Intelligence and Recombinant Growth," Working Paper (National Bureau of Economic Research, April 2018), <https://doi.org/10.3386/w24541>.
2. In order to understand the state of AI demand and adoption in Canada, Deloitte surveyed over 1,000 Canadian citizens and 2,500 businesses from around the world between July and September 2018. To understand early adopters, we surveyed leaders of 300 Canadian and 1,600 international large companies who are aggressively adopting AI to understand the practices and attitudes of early adopters. We asked them about their goals, spending, and results from adopting AI technologies, as well as the risks and challenges they perceive in implementing them.

All respondents were required to be knowledgeable about their company's use of artificial intelligence, and have direct involvement with their company's AI strategy, spending, implementation, and/or decision-making. Companies must report global annual revenues of at least \$50 million US and have at least 500 employees worldwide. The margin of error on these results is +/- 5.7 percentage points, 19 times out of 20.

To understand Canadian consumers' understanding, expectations, and beliefs about AI, we surveyed 1,019 Canadians from across the country. All respondents were aged 18+ and meant to be representative of the general adult population. The margin of error on these results is +/- 3.1 percentage points, 19 times out of 20.
3. Deloitte Canada, Citizen Survey 2018. See note above.
4. Valentina A. Assenova, "Modeling the Diffusion of Complex Innovations as a Process of Opinion Formation through Social Networks," *PLoS ONE* 13, no. 5 (May 2, 2018), <https://doi.org/10.1371/journal.pone.0196699>.
5. "Artificial Intelligence: A Policy-Oriented Introduction | Wilson Center," accessed March 1, 2019, <https://www.wilsoncenter.org/publication/artificial-intelligence-policy-oriented-introduction>.
6. "AI and Risk Management," Deloitte United States, accessed March 1, 2019, <https://www2.deloitte.com/global/en/pages/financial-services/articles/gx-ai-and-risk-management.html>.
7. Ivan Semeniuk, "What Is the Science behind the Cambridge Analytica Scandal?" *The Globe and Mail*, March 28, 2019 edition, accessed February 21, 2019, <https://www.theglobeandmail.com/world/article-what-is-the-science-behind-cambridge-analytica-scandal/>.
8. Zou, Bo, Feng Guo, and Jinyu Guo. "Absorptive Capacity, Technological Innovation, and Product Life Cycle: A System Dynamics Model." *SpringerPlus* 5, no. 1 (September 26, 2016). <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5037108/#CR46>.
9. Dietvorst, Simmons, and Massey, "Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err."
10. "Wielding Rocks and Knives, Arizonans Attack Self-Driving Cars - The New York Times," accessed February 20, 2019, <https://www.nytimes.com/2018/12/31/us/waymo-self-driving-cars-arizona-attacks.html>.
11. "AI Trust and AI Fears: A Media Debate That Could Divide Society — Oxford Internet Institute," accessed February 20, 2019, <https://www.oii.ox.ac.uk/blog/ai-trust-and-ai-fears-a-media-debate-that-could-divide-society/>.
12. Stephen Cave and Seán S. ÓhÉigeartaigh, "Bridging Near- and Long-Term Concerns about AI," *Nature Machine Intelligence* 1, no. 1 (January 2019): 5, <https://doi.org/10.1038/s42256-018-0003-2>.
13. Cave and ÓhÉigeartaigh.
14. "The Real Story of How the Internet Became so Vulnerable | The Washington Post," accessed February 20, 2019, <https://www.washingtonpost.com/sf/business/2015/05/30/net-of-insecurity-part-1/>.
15. "Draft Standards – CIO Strategy Council," accessed February 21, 2019, <https://ciostrategyCouncil.com/standards/draft-standards/>.
16. Janosch Delcker, "Finland's Grand AI Experiment," *POLITICO*, January 2, 2019, <https://www.politico.eu/article/finland-one-percent-ai-artificial-intelligence-courses-learning-training/>.
17. "Regulating Emerging Technology | Deloitte Insights," accessed February 21, 2019, <https://www2.deloitte.com/insights/us/en/industry/public-sector/future-of-regulation/regulating-emerging-technology.html>.
18. "The AI Revolution Needs a Rulebook. Here's a Beginning - The Globe and Mail," accessed December 13, 2018, <https://www.theglobeandmail.com/business/commentary/article-the-ai-revolution-needs-a-rulebook-heres-a-beginning/>.
19. "Understand and Control the Personal Information That You Store with Apple," Apple Support, accessed February 21, 2019, <https://support.apple.com/en-ca/HT208501>.
20. "Frequently Asked Questions —," accessed February 21, 2019, <https://www.crppc-gccamp.ca/faq-eng>.
21. Marten Kaevats, "Estonia Considers a 'Kratt Law' to Legalise Artificial Intelligence (AI)," *Medium* (blog), September 25, 2017, <https://medium.com/e-residency-blog/estonia-starts-public-discussion-legalising-ai-166cb8e34596>.
22. School of Policy Studies, Queen's University, "The Future of Work: What Do We Do?" - Session 6 - Armchair Discussion, accessed February 21, 2019, <https://www.youtube.com/watch?v=j8qZ21Rdmc8&t=0s&index=18&list=PL57JjtmhwgXnx2dFhCHmAsGfTxI7efYwS>.
23. "Navigating through New Forms of Work | Deloitte Insights," accessed February 21, 2019, <https://www2.deloitte.com/insights/us/en/deloitte-review/issue-21/navigating-new-forms-of-work.html>.
24. "Lifelong Learning Helps People, Governments and Business. Why Don't We Do More of It?," *World Economic Forum*, accessed February 21, 2019, <https://www.weforum.org/agenda/2017/07/lifelong-learning-helps-governments-business/>.
25. David Johnston, Trust: Twenty Reliable Ways to Build a Better Country (McClelland & Stewart, 2018).



omnia**AI**

Canada at 175 refers to Deloitte's vision for the future of our nation's prosperity by July 1, 2042. Now is the time for Canada to redefine what it means to be a global leader.

Our ambition as a nation must be bolder—we must aim to not only achieve economic prosperity, but to elevate our quality of life for our vast and diverse population.

Deloitte believes that in the years leading up to our 175th anniversary Canada can remain in a class by itself—not one of the best, but the best place in the world to live and work. To do this, we must adopt an unrelenting commitment to developing courageous leaders who embody an inclusive way of being.

Contact us at canada175@deloitte.ca

We believe AI has the power to improve Canadian organizations in transformative ways—and we'll work with you to support your AI journey from strategy to scalable implementation.

At Omnia AI, we provide end-to-end solutions to complex challenges. We see new opportunities in emerging technologies that others don't. And because we're Deloitte, we're trusted advisors who work to understand your organization's objectives from all angles.

As AI is poised to rapidly change the world, we're here to:

- Show you AI's potential to tackle core organizational priorities—revenue, cost, risk, and experience— head on
- Demonstrate how AI can bring new value to your organization starting today, and set the foundation for growth tomorrow
- Bring the right people and advice to your challenges, along with the right technology to maximize the benefits

Connect with Omnia AI at www.deloitte.ca/OmniaAI

Deloitte.

www.canada175.ca

Deloitte provides audit and assurance, consulting, financial advisory, risk advisory, tax, and related services to public and private clients spanning multiple industries. Deloitte serves four out of five Fortune Global 500® companies through a globally connected network of member firms in more than 150 countries and territories bringing world-class capabilities, insights and service to address clients' most complex business challenges. To learn more about how Deloitte's approximately 264,000 professionals—9,400 of whom are based in Canada—make an impact that matters, please connect with us on LinkedIn, Twitter or Facebook.

Deloitte LLP, an Ontario limited liability partnership, is the Canadian member firm of Deloitte Touche Tohmatsu Limited. Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee, and its network of member firms, each of which is a legally separate and independent entity. Please see www.deloitte.com/about for a detailed description of the legal structure of Deloitte Touche Tohmatsu Limited and its member firms.

© Deloitte LLP and affiliated entities.

Designed and produced by the Deloitte Design Studio, Canada. 18-5963M